QuakeSim: Efficient Modeling of SensorWeb Data in a Web Services Environment

Andrea Donnellan

Jet Propulsion Laboratory,

California Institute of Technology

NASA Earth Science Technology Conference

June 25, 2008

Jay Parker
John Rundle
Robert Granat
Geoffrey Fox
Dennis McLeod
Marlon Pierce
Lisa Grant
Walter Brooks
Rami Al-Ghanmi

QuakeSim

- Modeling environment for studying earthquake processes using a web services environment
  - *Focuses on modeling interseismic processes*
- Federates data from multiple sources and integrates the databases with modeling applications
  - *Spaceborne GPS and InSAR data, geological fault data, and seismicity data*
- Applications include various boundary element, finite element, and analytic applications
  - *Run on a range of platforms including desktop and high end computers*
  - *Used to simulate interacting earthquake fault systems, model nucleation and slip on faults, and calculate run-up and inundation from tsunamis generated by offshore earthquakes*
  - *Applies pattern recognition techniques to real and simulated data to elucidate subtle features in the processes.*

QuakeSim Overview

- QuakeSim merges three approaches
  - *Modeling, simulation, and analysis tools*
  - *Web services portal environment*
  - *High performance computing*

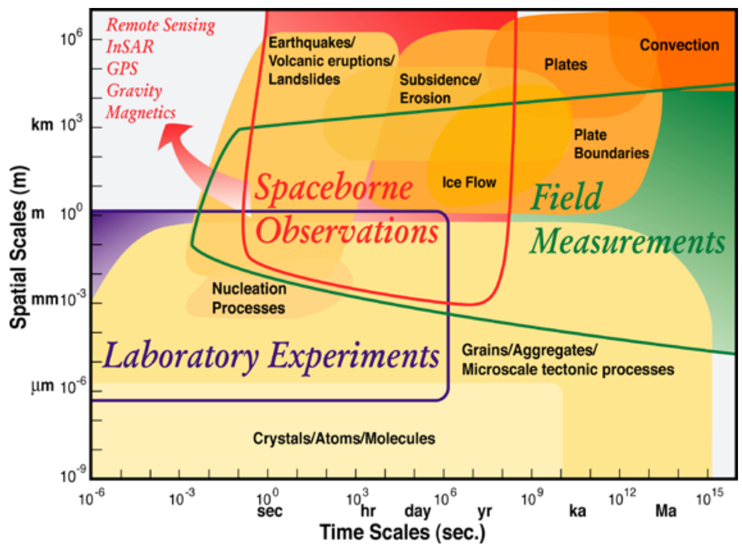

- Focuses on the interseismic part of the earthquake cycle

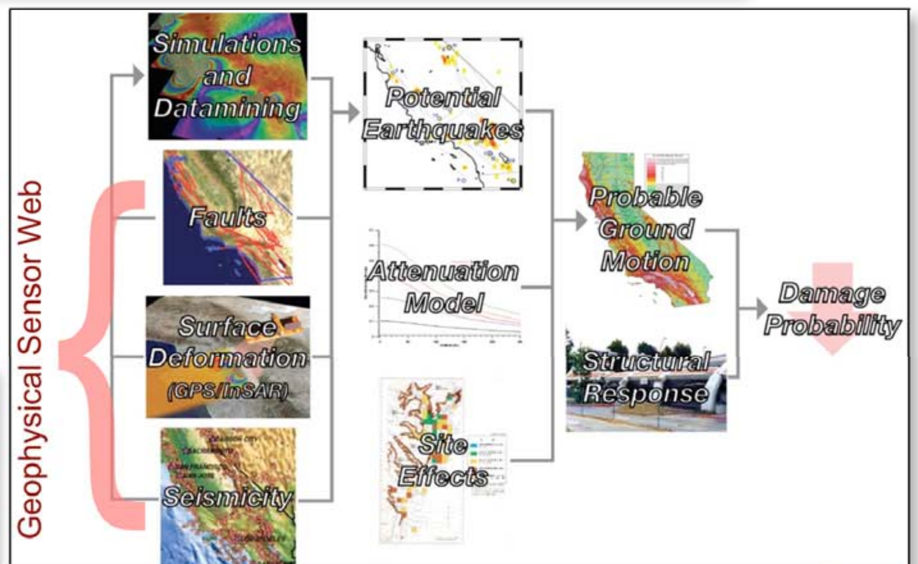

- Developed with an interest toward ingesting spaceborne deformation data
  - *GPS*
  - *InSAR (DESDynI)*

Portlet Summary

RDAHMM	Set up and run RDAHMM, query Scripps GRWS GPS Service, maintain persistent user sessions
ST_Filter	Similar to RDAHMM portlet; ST_Filter has much more input
Station Monitor	Shows GPS stations on a Google Map, displays last 10 minutes of data
Real Time RDAHMM	Displays RDAHMM results of last 10 minutes of GPS data in a Google map
Daily RDAHMM	Calculates, updates RDAHMM event classifications with daily updated GPS data from SOPAC's GRWS service (14 day delay, but uses all the data)
GeoFEST	Create input geometries, generate FE meshes, run parallel FEM solvers
Disloc, Simplex	Calculate service displacements from fault models

Grid Services

- Have implemented GoeFEST through the portal on the NSF TeraGrid
  - *GRAM servers at TACC's (Texas Advanced Computing Center) LoneStar; NCSA's (National Center for Supercomputer Applications) Cobalt, Tungsten, and Mercury; and clusters at Oak Ridge and University of Chicago*
- Cosmos at JPL
  - *Using the JPL extranet with globus to access the machines*
- Columbia at NASA Ames
  - *QuakeSim represents the first major user of Columbia that has identified Grid services and Condor G as a requirement for job launch*

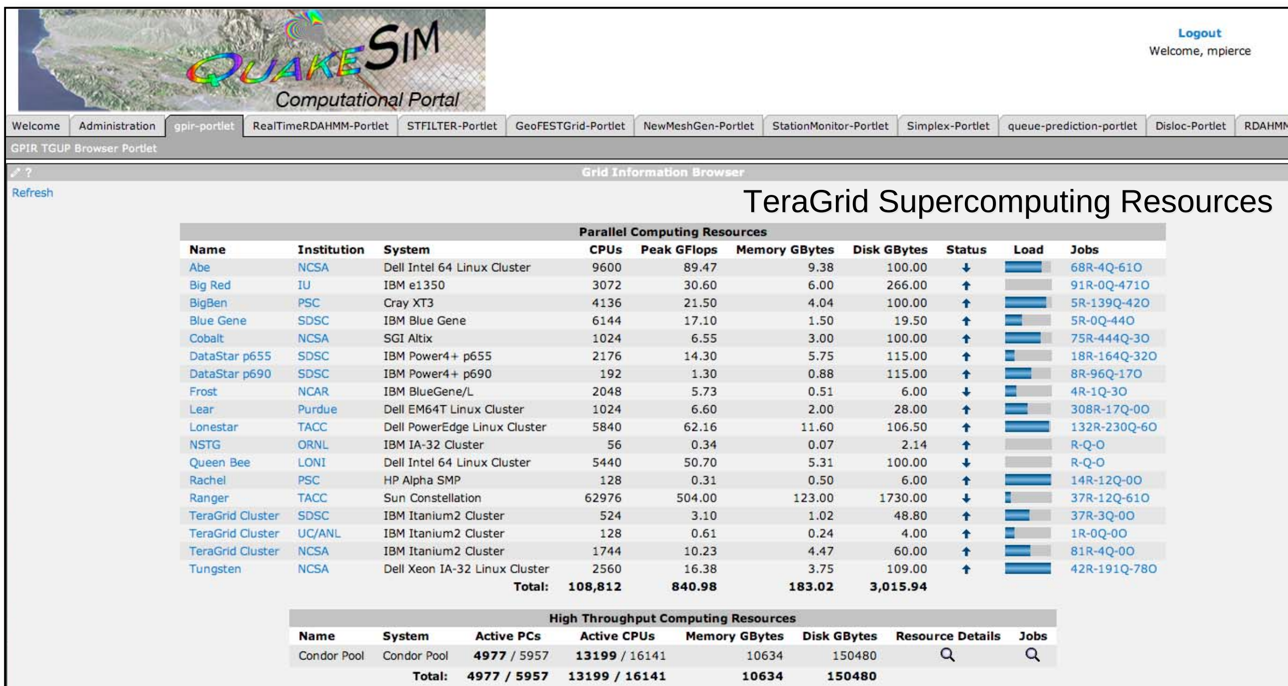

Parallel Computing Resources

Name	Institution	System	CPUs	Peak GFlops	Memory GBytes	Disk GBytes	Status	Load	Jobs
Abe	NCSA	Dell Intel 64 Linux Cluster	9600	89.47	9.38	100.00	↓		68R-4Q-610
Big Red	IU	IBM e1350	3072	30.60	6.00	266.00	↑		91R-0Q-4710
BigBen	PSC	Cray XT3	4136	21.50	4.04	100.00	↑		5R-139Q-420
Blue Gene	SDSC	IBM Blue Gene	6144	17.10	1.50	19.50	↑		5R-0Q-440
Cobalt	NCSA	SGI Altix	1024	6.55	3.00	100.00	↑		75R-444Q-30
DataStar p655	SDSC	IBM Power4+ p655	2176	14.30	5.75	115.00	↑		18R-164Q-320
DataStar p690	SDSC	IBM Power4+ p690	192	1.30	0.88	115.00	↑		8R-96Q-170
Frost	NCAR	IBM BlueGene/L	2048	5.73	0.51	6.00	↓		4R-1Q-30
Lear	Purdue	Dell EM64T Linux Cluster	1024	6.60	2.00	28.00	↑		308R-17Q-00
Lonestar	TACC	Dell PowerEdge Linux Cluster	5840	62.16	11.60	106.50	↑		132R-230Q-60
NSTG	ORNL	IBM IA-32 Cluster	56	0.34	0.07	2.14	↑		R-Q-0
Queen Bee	LONI	Dell Intel 64 Linux Cluster	5440	50.70	5.31	100.00	↓		R-Q-0
Rachel	PSC	HP Alpha SMP	128	0.31	0.50	6.00	↑		14R-12Q-00
Ranger	TACC	Sun Constellation	62976	504.00	123.00	1730.00	↓		37R-12Q-610
TeraGrid Cluster	SDSC	IBM Itanium2 Cluster	524	3.10	1.02	48.80	↑		37R-3Q-00
TeraGrid Cluster	UC/ANL	IBM Itanium2 Cluster	128	0.61	0.24	4.00	↑		1R-0Q-00
TeraGrid Cluster	NCSA	IBM Itanium2 Cluster	1744	10.23	4.47	60.00	↑		81R-4Q-00
Tungsten	NCSA	Dell Xeon IA-32 Linux Cluster	2560	16.38	3.75	109.00	↑		42R-191Q-780
		Total:	**108,812**	**840.98**	**183.02**	**3,015.94**			

High Throughput Computing Resources

Name	System	Active PCs	Active CPUs	Memory GBytes	Disk GBytes	Resource Details	Jobs
Condor Pool	Condor Pool	**4977** / 5957	**13199** / 16141	10634	150480		
	Total:	**4977 / 5957**	**13199 / 16141**	**10634**	**150480**		

Globus Support and Upgrading of the Portal

- Build portals out of portlets (Java Standard)
  - *Reuse capabilities from our Open Grid Computing Environments (OGCE) project, the REASoN GPS Explorer project, and many TeraGrid Science Gateways*
  - *Decorate with Google Maps, Yahoo UI gadgets, etc.*
- Use Java Server Faces to build individual component portlets
  - *Build standalone tools, then convert to portlets at the very end*
- Use simple Web Services for accessing codes and data
  - *Keep It Stateless …*
- Use Globus job and file management services for interacting with high performance computers
- Favor Google Maps and Google Earth for their simplicity, interactivity and open APIs
  - *Generate KML and GeoRSS*
- Use Apache Maven based build and compile system

Web Services Upgrade

- The current version of the portal is much more stable than the previous version
  - *On a related project hundreds of simultaneous users were successfully simulated*
- The portlets also don't have some state bugs that the older jetspeed had
  - *Previously we had a lot of code that tried to catch everything to keep from, for example, resubmitting jobs or loading null variables when one flipped around through the tabs*
  - *The current container and JSF portlet bridge we are using manages all of this without our worrying about it*
- Also re-factored a lot of the code to move things out of the portal and into services where they belong
  - *This makes the services a lot more independent of the portal, which has some value if we ever want to integrate with a workflow engine*
- Much of the original code is preserved, it is just moved around

Upgrading Portal

QuakeSim, Version 1	Reason to Revise	QuakeSim, Version 2
Application Web Service for wrapping a.out executables. Execution management service built with Apache Ant.	Services **too coupled to portal**; no simple WSDL programming interface; could not be used in workflow engines; not self contained	Give each code a **proper service interfac**e. Retain Apache Ant core but extend. Keep WSDL message structure simple (Strings, ints, doubles, URLs), wrapped as Java Beans
File Management Service	Unnecessary, too coupled to Apache Axis 1.0	HTTP GET, URLs
Context Management Service manages persistent portal sessions using recursive XML structure.	**Too slow** (file system); didn't scale; XML databases didn't mature; Object-Relational Mappings (ORM) not efficient	Using **DB40**; all services communicate with easily XML serializable JavaBeans
OGC-compatible map and data services	**Too complicated**; ORM is a big overhead.	Google Maps, KML generating services
Serial job submission	NSF TeraGrid and Open Science Grid run full time **production Grids** for HPC	**Condor-G/Birdbath** based job management extensions to GeoFEST service

Earthquake Forecasting

- In the last six months five earthquakes above magnitude 5 have occurred in identified hotspots
- The identified hotspots make up only 1.2% of the total map area of the forecast or the state of California
  - *Approximately half the total boxes have at least 1 M>3 earthquake in them*
  - *Therefore for the total active area, it would be forecast area of 2.4%*
- The approach is to minimize the forecast area, which is essentially the false alarm rate, while still detecting all the large earthquakes (maximizing the hit rate)
- The mean forecast error is the average distance that a M>5 earthquake occurs from a 11 km pixel box boundary
  - *Most of this error is due to offshore earthquake # 6, which is about 50 km off the nearest red pixel*
  - *Without that earthquake, the error would be less than 5.5 km, half of one pixel box size*
- JR Holliday, CC Chen, KF Tiampo, JB Rundle, DL Turcotte and A Donnellan, "A RELM Earthquake Forecast Based on Pattern Informatics," Seism. Res. Lett., v78, pp. 87-93 (2007)

Earthquake Forecasting: 2002 Scorecard

Status of the Real Time Earthquake Forecast Experiment (Original Version)

(JB Rundle et al., PNAS, v99, Supl 1, 2514-2521, Feb 19, 2002; KF Tiampo et al., Europhys. Lett., 60, 481-487, 2002; JB Rundle et al.,Rev. Geophys. Space Phys., 41(4), DOI 10.1029/2003RG000135,2003. http://quakesim.jpl.nasa.gov)

How are We Doing?
(Composite N-S Catalog)

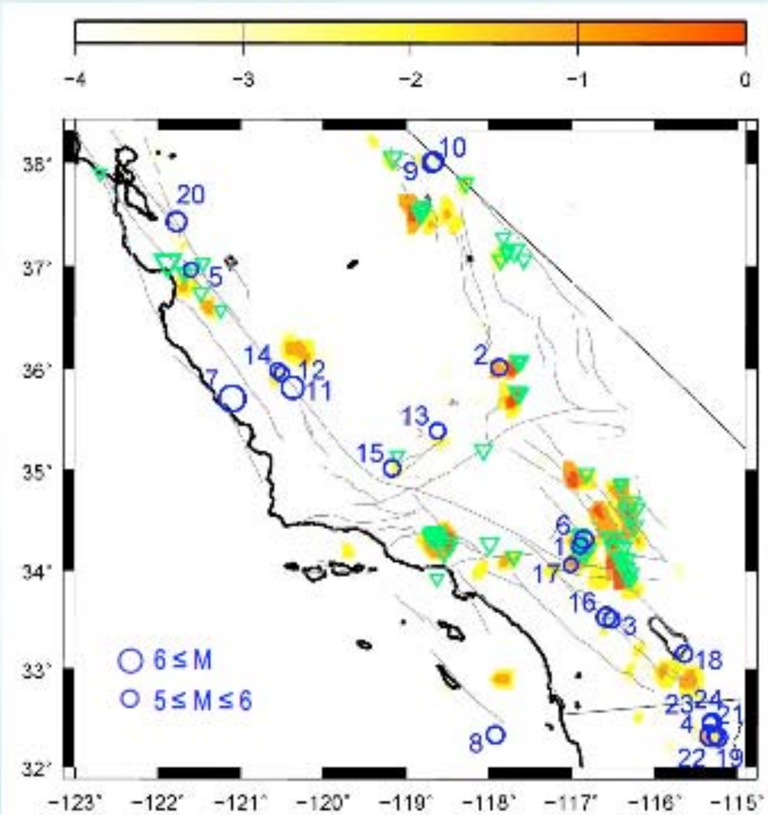

Plot of Log_{10} (Seismic Potential)

Increase in Potential for significant earthquakes, ~ 2000 to 2010

Twenty- four significant earthquakes (blue circles) have occurred in Central or Southern California. Margin of error of the anomalies is +/- 11 km; Data from S. CA. and N. CA catalogs:

After the work was completed

1. Big Bear I, M = 5.1, Feb 10, 2001
2. Coso, M = 5.1, July 17, 2001

After the paper was in press (September 1, 2001)

3. Anza I, M = 5.1, Oct 31, 2001

After the paper was published (February 19, 2002)

4. Baja I, M = 5.7, Feb 22, 2002
5. Gilroy, M=4.9 - 5.1, May 13, 2002
6. Big Bear II, M=5.4, Feb 22, 2003
7. San Simeon, M = 6.5, Dec 22, 2003
8. San Clemente Island, M = 5.2, June 15, 2004
9. Bodie I, M=5.5, Sept. 18, 2004
10. Bodie II, M=5.4, Sept. 18, 2004
11. Parkfield I, M = 6.0, Sept. 28, 2004
12. Parkfield II, M = 5.2, Sept. 29, 2004
13. Arvin, M = 5.0, Sept. 29, 2004
14. Parkfield III, M = 5.0, Sept. 30, 2004
15. Wheeler Ridge, M = 5.2, April 16, 2005
16. Anza II, M = 5.2, June 12, 2005
17. Yucaipa, M = 4.9 - 5.2, June 16, 2005
18. Obsidian Butte, M = 5.1, Sept. 2, 2005
19. Baja II, M = 5.4, May 23, 2006
20. Alum Rock, M=5.6, Oct. 30, 2007
21. Baja III, M = 5.4, Feb 9, 2008
22. Baja IV, M = 5.1, Feb 11, 2008
23. Baja V, M = 5.0, Feb 12, 2008
24. Baja VI, M=5.0, Feb 19, 2008

Note: This original forecast was made using both the full Southern California catalog plus the full Northern California catalog. The S. Calif catalog was used south of lattitude 36°, and the N. Calif. catalog was used north of 36°. No corrections were applied for the different event statistics in the two catalogs. Green triangles mark locations of large earthquakes (M ≥ 5.0) between Jan 1, 1990 – Dec 31, 1999.

CL#03-2015

Earthquake Forecasting: 2007 Scorecard

New California Earthquake Forecast Scorecard

Map published in: JR Holliday et al., Seism. Res. Lett., v 78, Jan/Feb 2007 pp. 87-93

Red pixels (higher risk locations) are based on data up to August 31, 2005

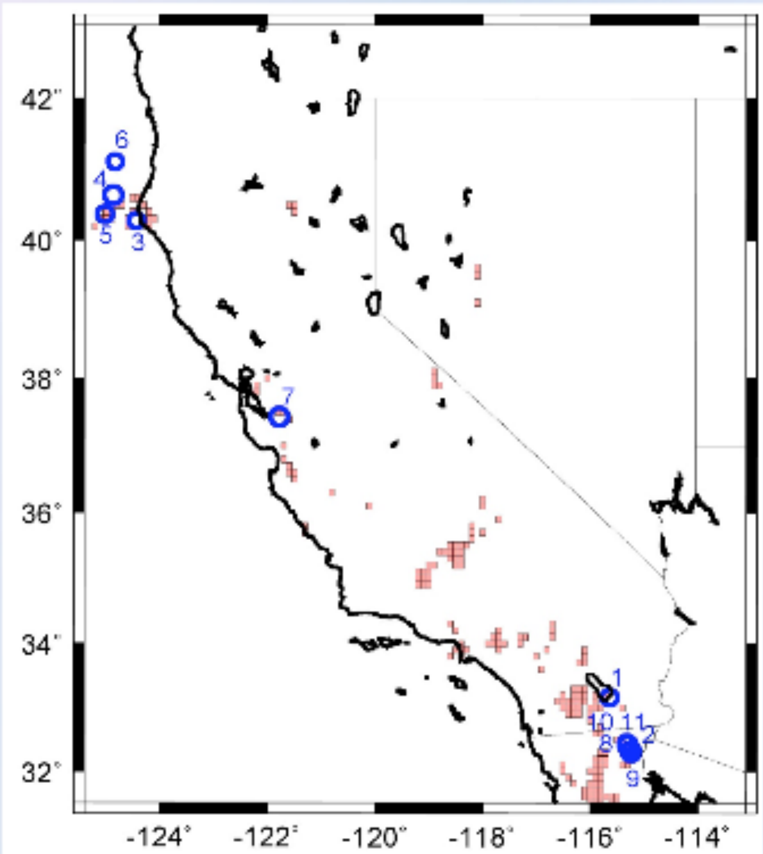

Blue circles are earthquakes M > 5.0 occurring after September 1, 2005 up to the present

(Dates and times are UTC)

1. Sept 2, 2005, M = 5.1
2. May 24, 2006, M = 5.4
3. July 19, 2006, M = 5.0
4. Feb 26, 2007, M = 5.4
5. May 9, 2007, M = 5.2
6. June 25, 2007, M = 5.0
7. Oct 30, 2007, M = 5.5
8. Feb 9, 2008, M = 5.4
9. Feb 11, 2008, M = 5.1
10. Feb 12, 2008, M = 5.0
11. Feb 19, 2008, M = 5.0

Mean Forecast Error = 6 ± 15 km
@ 1.2% forecast area coverage

For M > 6.0 EQs, N. California is now at HIGHER Risk △
For M > 6.0 EQs, S. California is now at LOWER Risk ▽

Temporal risk is calculated using the method published in JR Holliday et al, Phys. Rev. Lett., v97, p 238 (2006)

Earthquakes change the state of the system

Therefore - want to update the forecast continuously

However - need to test the methodology, hence static forecasts

Virtual California: Interacting Fault Systems

- Select faults – edit fault topology
- Select friction model – edit friction model
- Run earthquake simulation code under user – defined conditions
- Compute surface deformation corresponding to slip history
- Visualize – InSAR fringes & other items, including fault topology, friction models, slip distributions, etc.
- Analyze – statistics toolbox
- Assimilation protocols – adjust model parameters to optimize, based on paleoseismic and other data

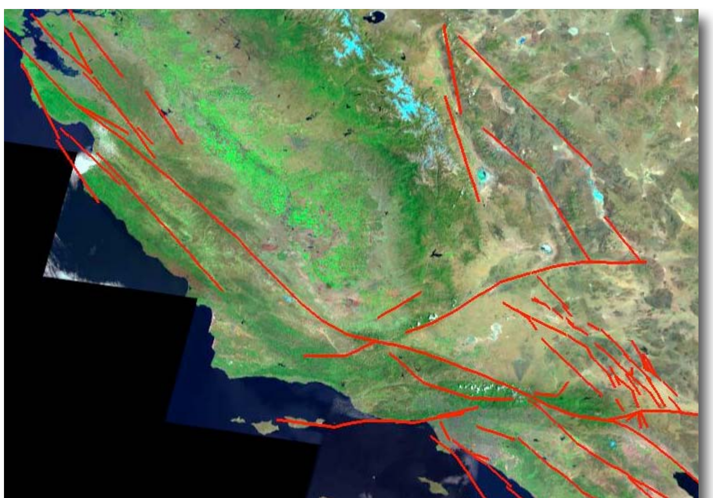

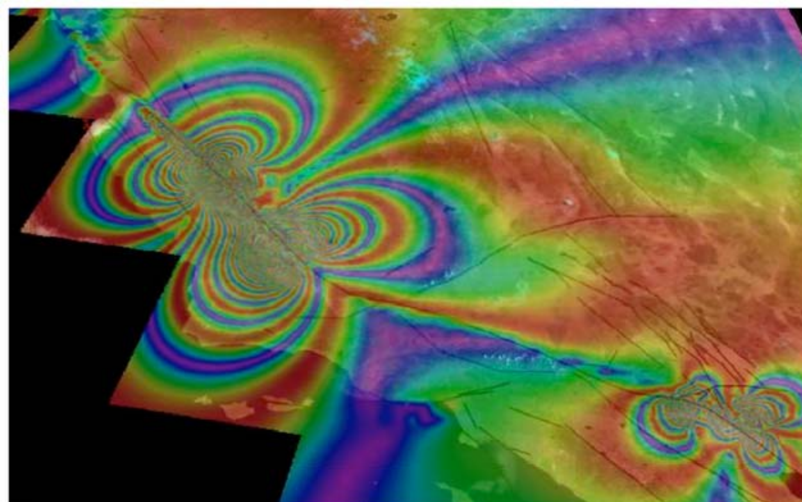

Statistical Analysis of Virtual California Data

Systematic observations over a long time period provide an opportunity to observe emergent behavior and fault interactions in the system

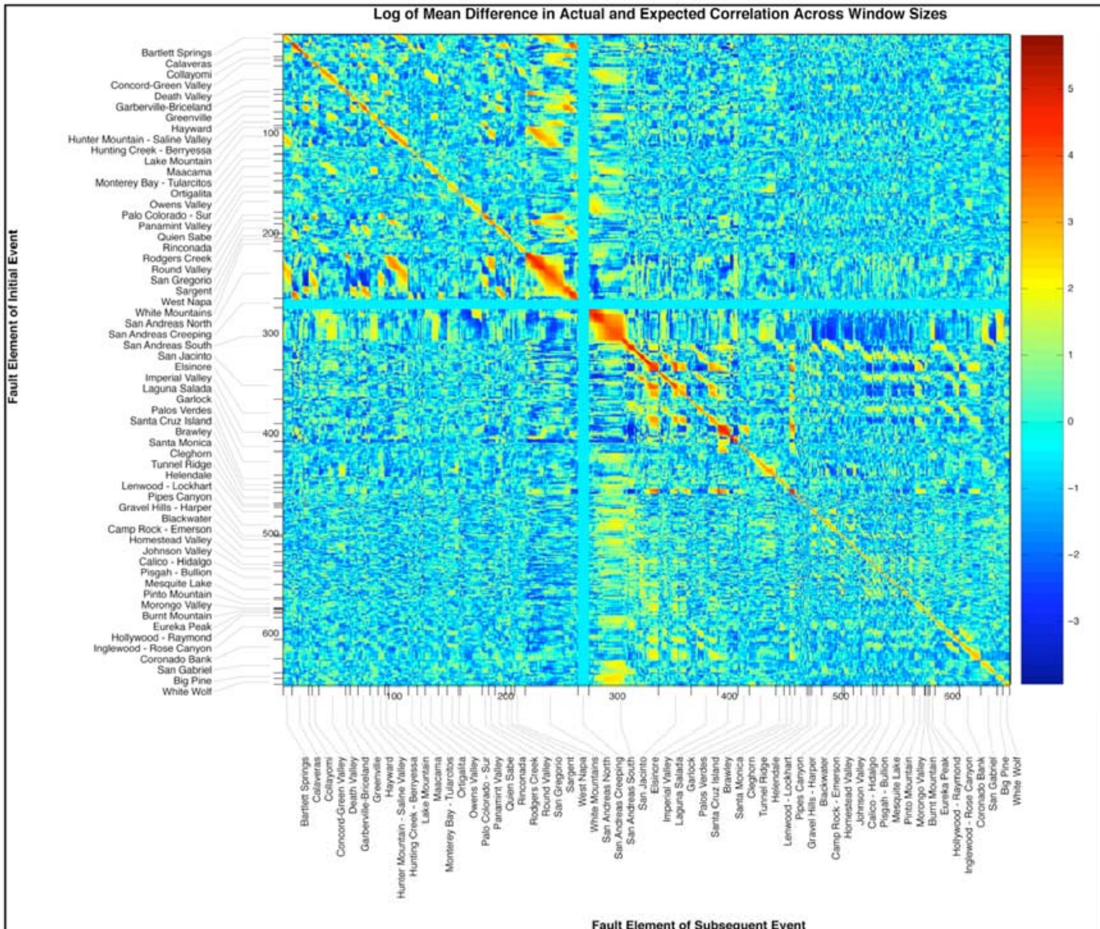

- Virtual California produces a very large synthetic seismic record
- Calculate the correlation of events on an “initiating” fault segment with subsequent events on a second fault element
- Produces correlation score matrix that shows the relative amount of correlation between events on two elements
- Provides insight into the relationships between faults at different time scales
- Analysis shows that events on the southern San Andreas fault typically follow but do not precede events on the Eastern California Shear Zone

Log of correlation score matrix 400 yr time window. Log function is applied to highlight interesting features. Does not include the “creeping” section of the San Andreas fault.

RDAHMM Time Series Analysis

- RDAHMM: Regularized Deterministic Annealing Hidden Markov Model
  - *Identifies state changes in GPS time series data*
    - Aquifer water withdrawal and associated subsidence
    - Earthquakes
    - GPS data processing reference frame changes
  - *Are now producing daily updates*

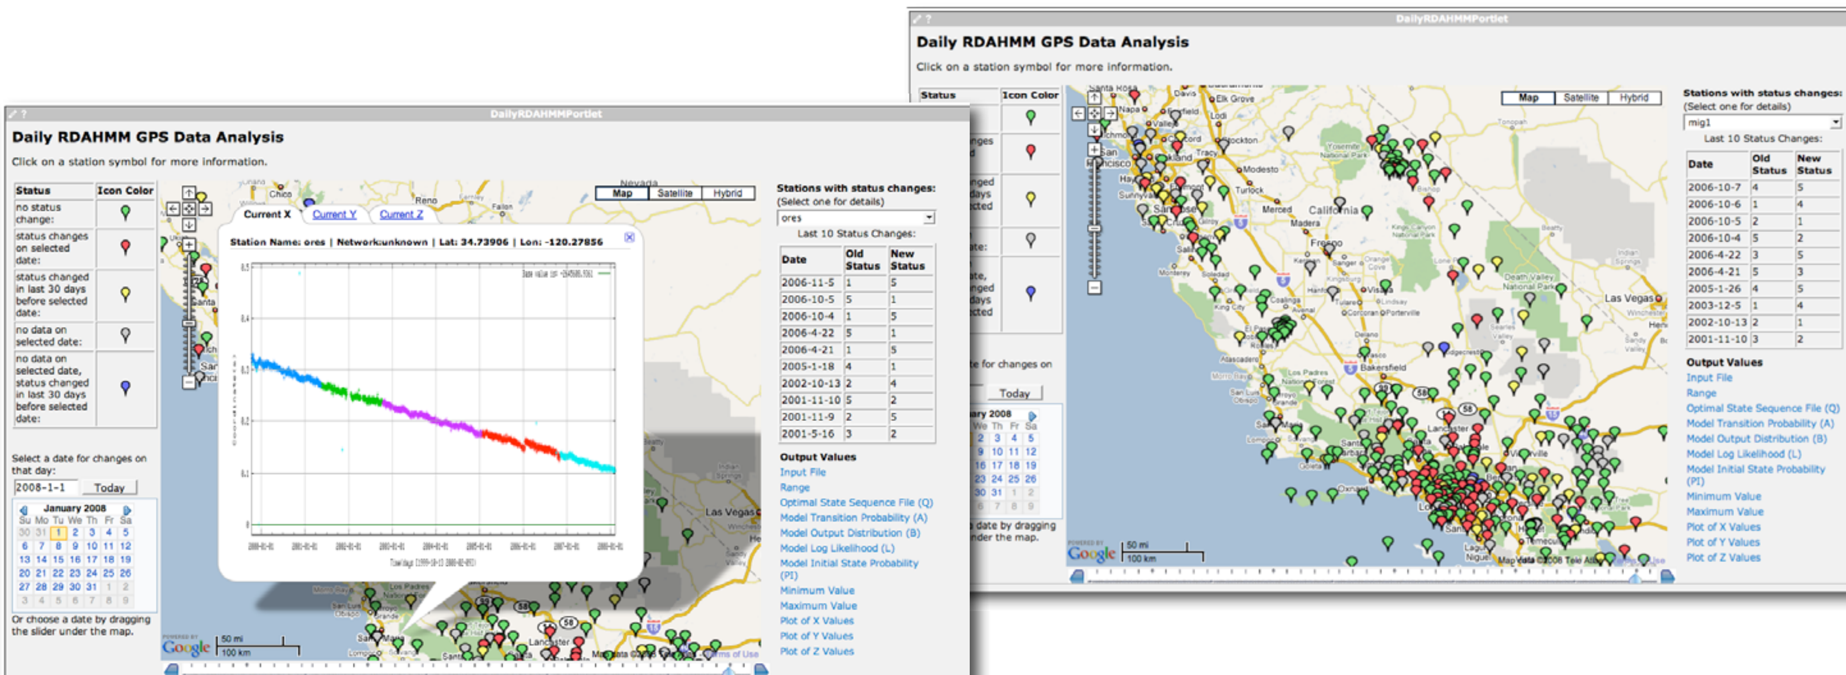

Using RDAHMM to Identify State Changes

- RDAHMM identified a state change of many stations in SCIGN (Southern California Integrated GPS Network)
- Investigation indicates that a reference frame change due to an earthquake in Siberia that day is the most likely cause of the state change
- The impact to the network position time series data would not have been identified without RDAHMM

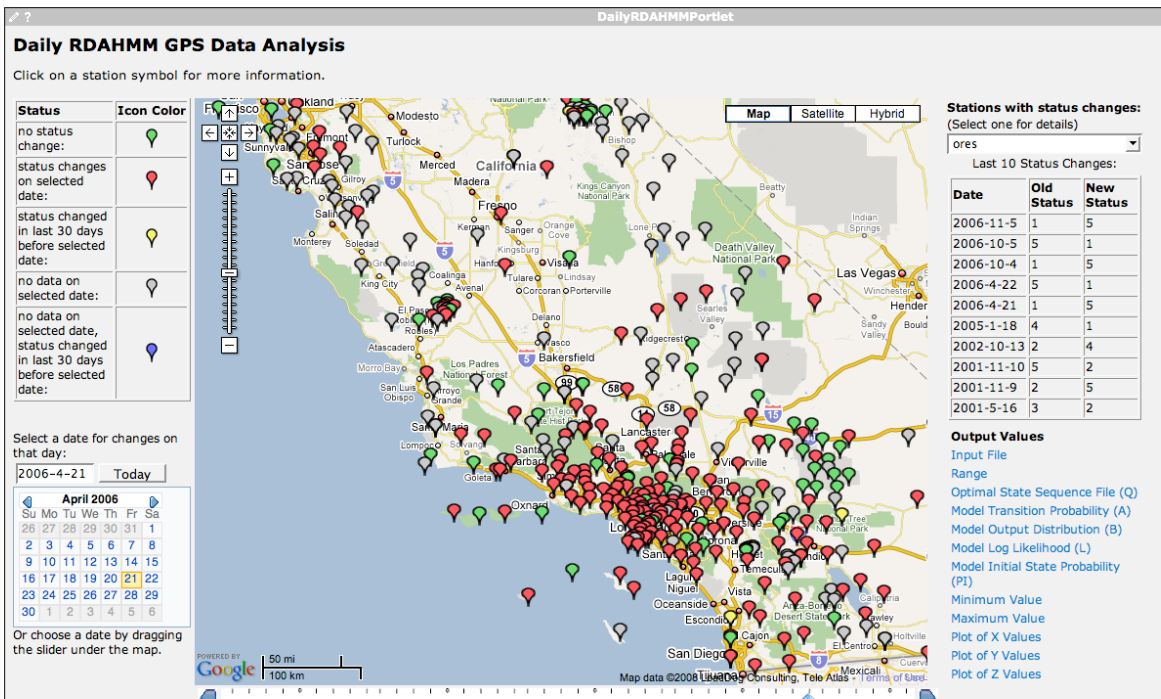

1906 San Francisco Earthquake Models

- Investigating the effects of large earthquakes over time
- Large-scale simulations are using portal tools, QuakeSim codes, and supercomputing time on Project Columbia in order to produce model results that enhance our understanding of the earthquake process
- More complex (realistic) models show postseismic effects of the 1906 earthquake of 1 mm/yr rather than the 2–5 mm/yr for the less complex models

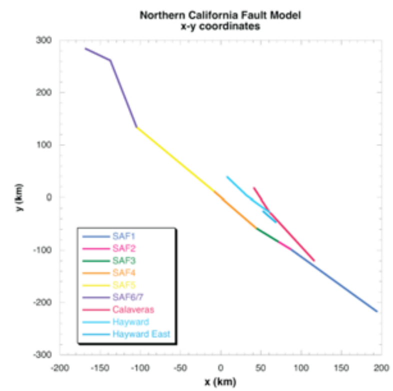

Multiple fault model segment geometry

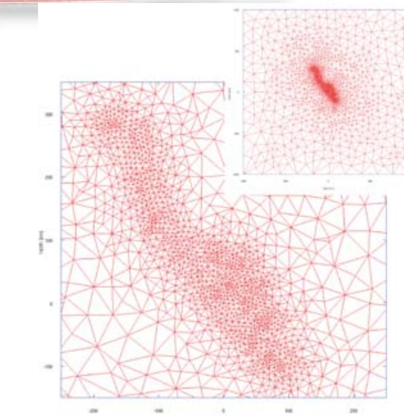

Multiple fault model mesh

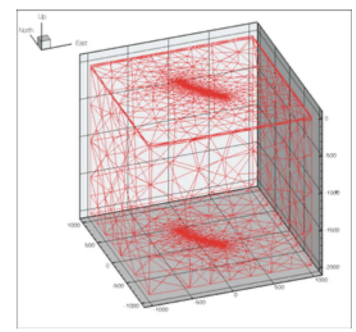

Multiple segment SAF model mesh

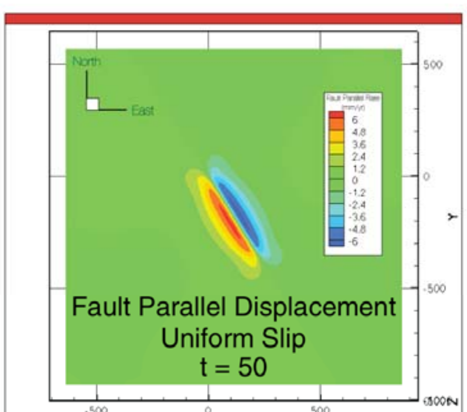

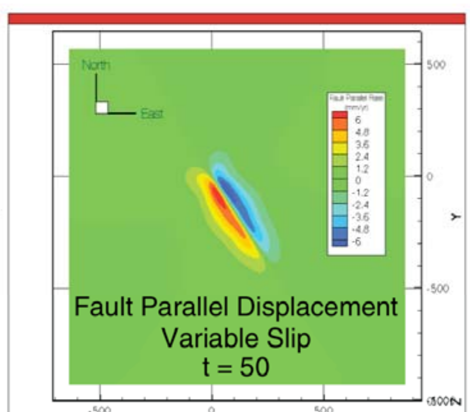

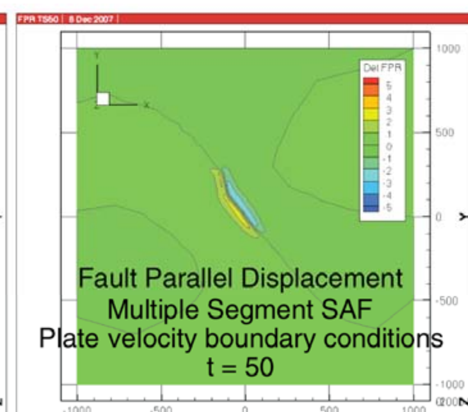

QuakeSim Federated Database System Architecture

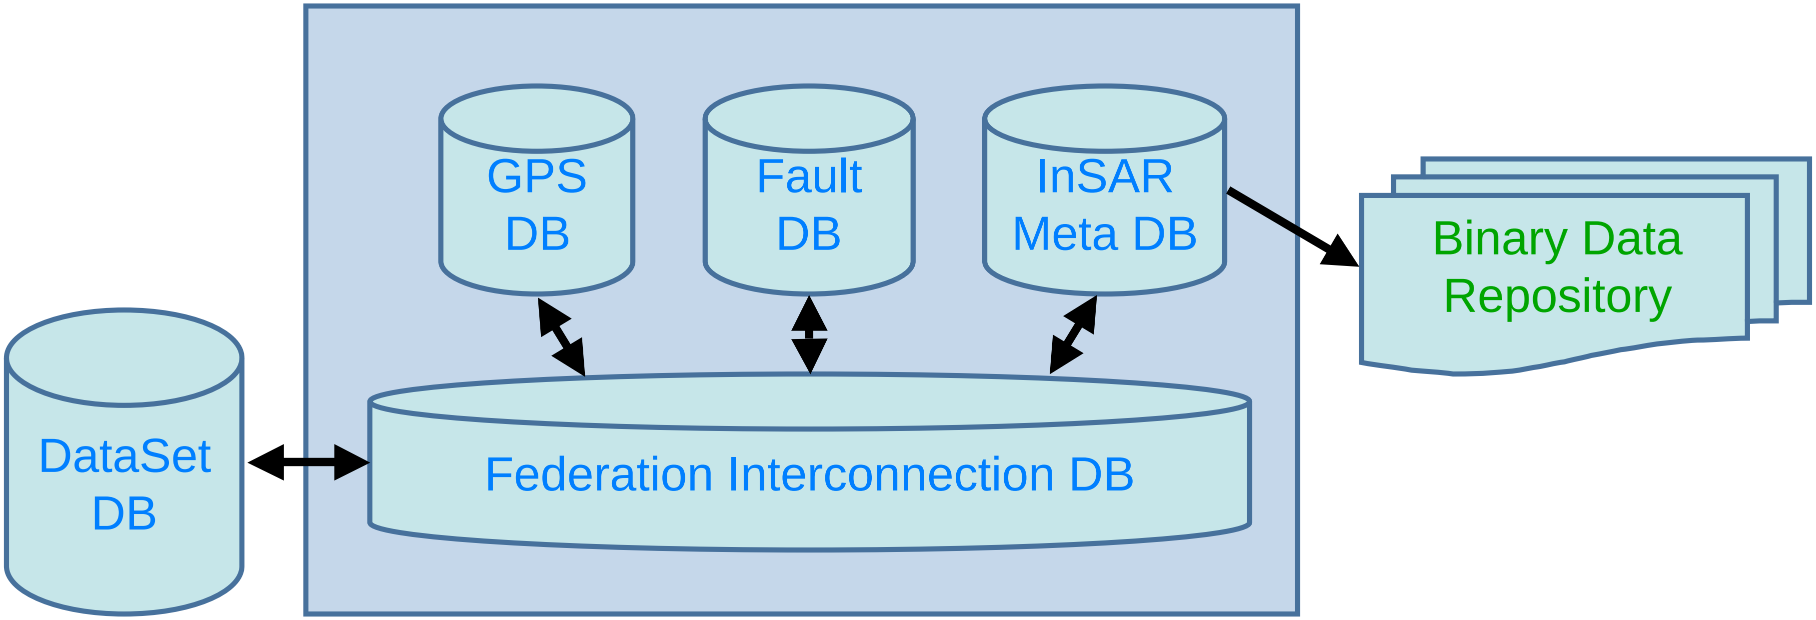

- Databases specified at ontology, EER, and relational levels
- DataSet database contains source literature and documentation
- Binary Data Repository contain source InSAR data
- Interface via portal, data exchange via web services

InSAR Database

- Paul Lundgren at JPL has been processing every available scene in southern California and is generously making them available to QuakeSim
- We are currently populating the database
- Next step will be to access the database with Simplex (fault movement inversion application)

EER Specification

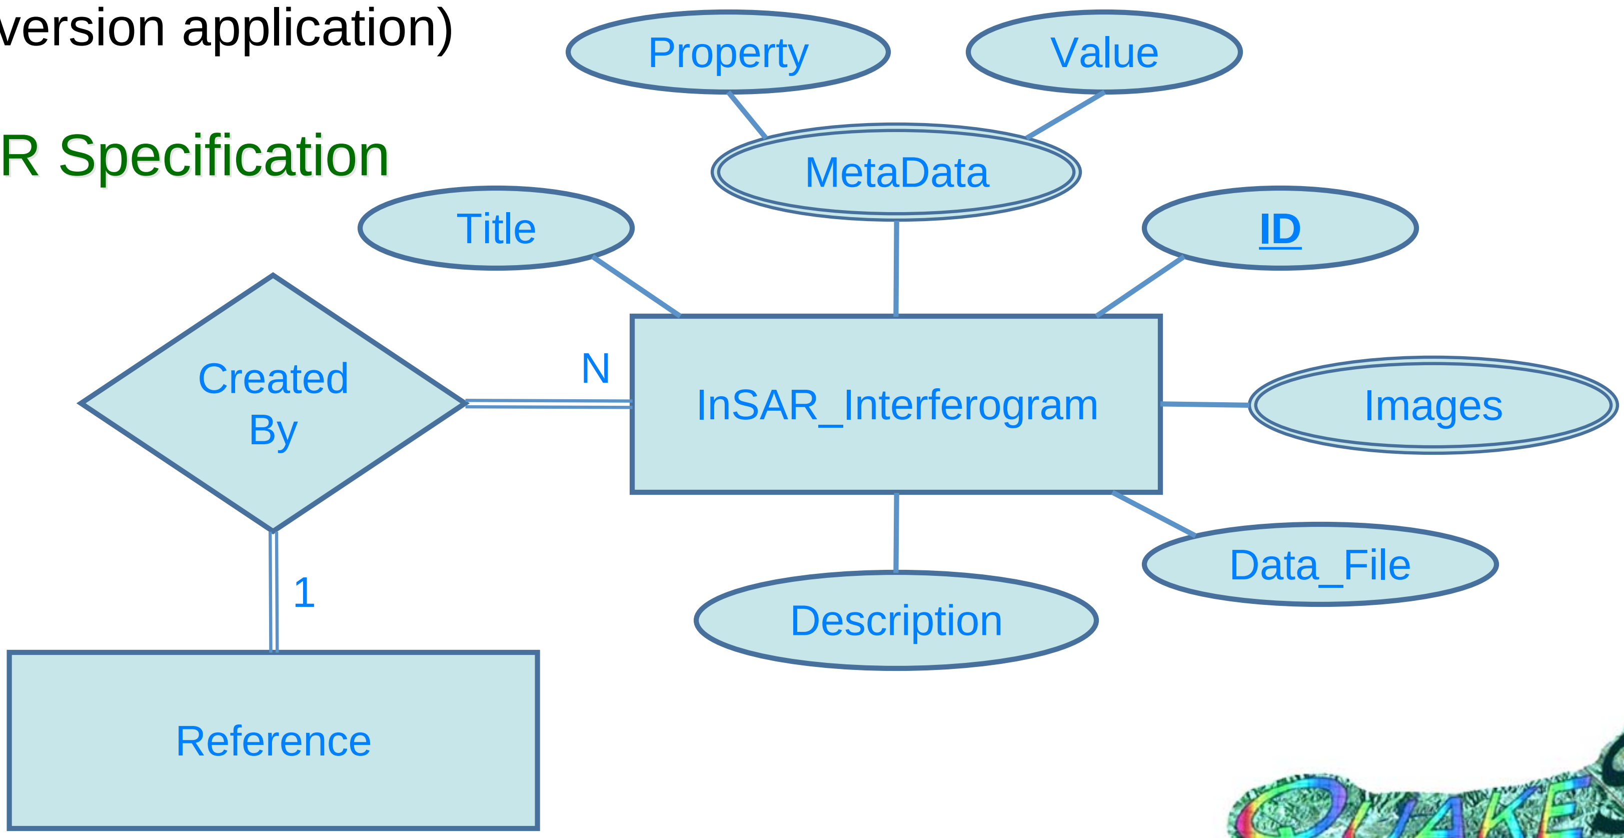

QuakeTables InSAR Database

Quake Tables

Jet Propulsion Laboratory
California Institute of Technology

QuakeSIM

Indiana University · Brown University · University of Southern California · UC Davis · UC Irvine

InSAR Database

Insert InSAR Data

Last modified on 02/20/2008

Thumbnail	Title	Description	Download
	Northridge, California 1992-06-01 to 1995-07-19	InSAR Image for Northridge, California by Paul Lundgren	metadata - data
	Northridge, California 1992-09-14 to 1995-11-02	InSAR Image for Northridge, California by Paul Lundgren	metadata - data
	Northridge, California 1992-11-23 to 1996-01-10	InSAR Image for Northridge, California by Paul Lundgren	metadata - data
	Northridge, California 1992-11-23 to 1996-07-04	InSAR Image for Northridge, California by Paul Lundgren	metadata - data
	Northridge, California 1993-08-30 to 1995-04-05	InSAR Image for Northridge, California by Paul Lundgren	metadata - data
	Northridge, California 1993-11-08 to 1995-12-06	InSAR Image for Northridge, California by Paul Lundgren	metadata - data

Browsing the Database

Interferogram

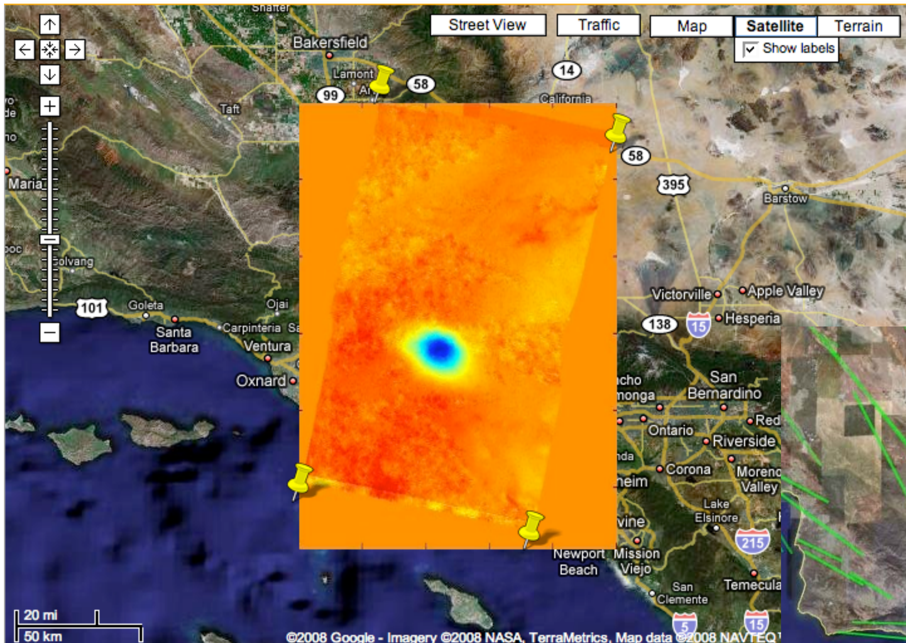

Fault Data

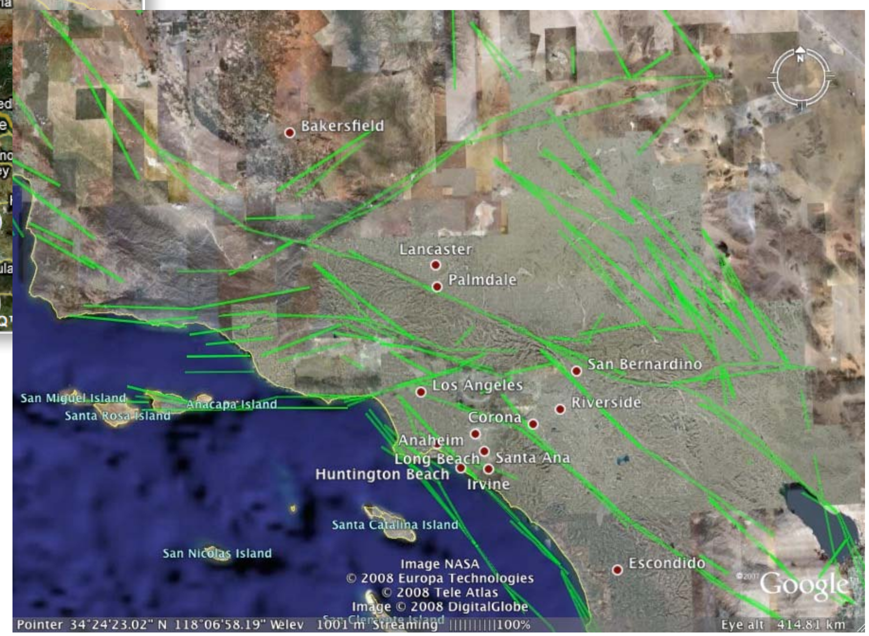

DESDynI Mission

Deformation, Ecosystem Structure, and Dynamics of Ice

- Two sensors to provide observations for
  - *Solid Earth (surface deformation)*
  - *Ecosystems (terrestrial vegetation structure)*
  - *Climate (ice dynamics)*
- Five year mission
- Frequent (8-day) revisit
- L-band synthetic aperture radar (SAR) system
  - *Operated as a repeat-pass interferometer (InSAR)*
  - *Multiple polarization: single, dual, or fully polarimetric*
  - *Strip-map or scanSAR modes with a viewable swath of 340 km*
  - *35 m ground resolution*
  - *Two sub-bands for ionospheric correction*
- Multiple-beam lidar
  - *Operating in the infrared (1,064 nm)*
  - *25-m spatial resolution*
  - *Canopy-height accuracy of 1 m*

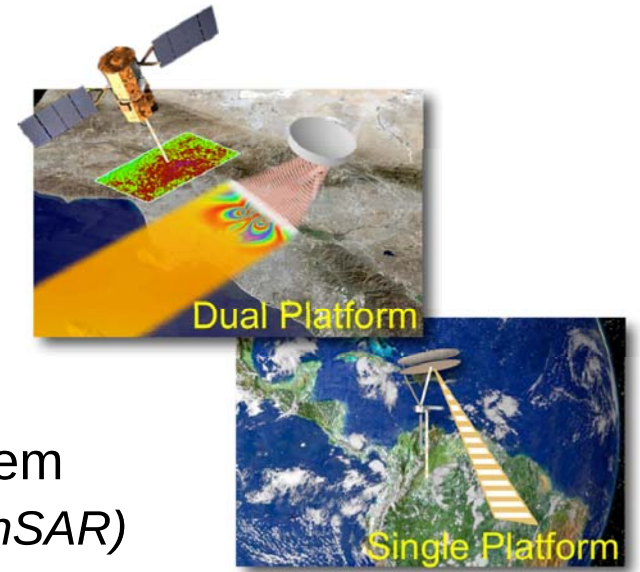

Observation Techniques

- DESDynI is defined by the decadal survey as an **L-band InSAR** and **multibeam LIDAR** mission for improving our understanding of hazards, ice sheet dynamics, and ecosystems

Repeat Pass InSAR

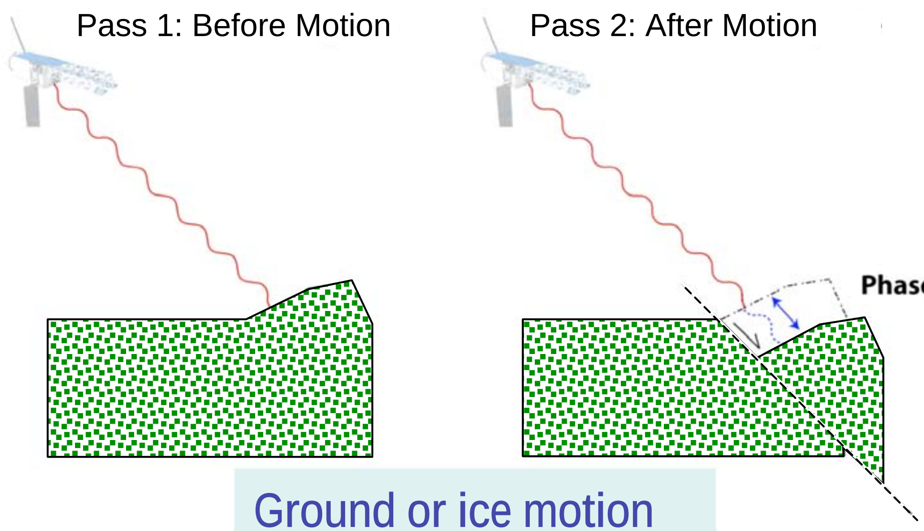

Ground or ice motion

Polarimetric SAR and Finite Baseline InSAR

Multibeam LIDAR

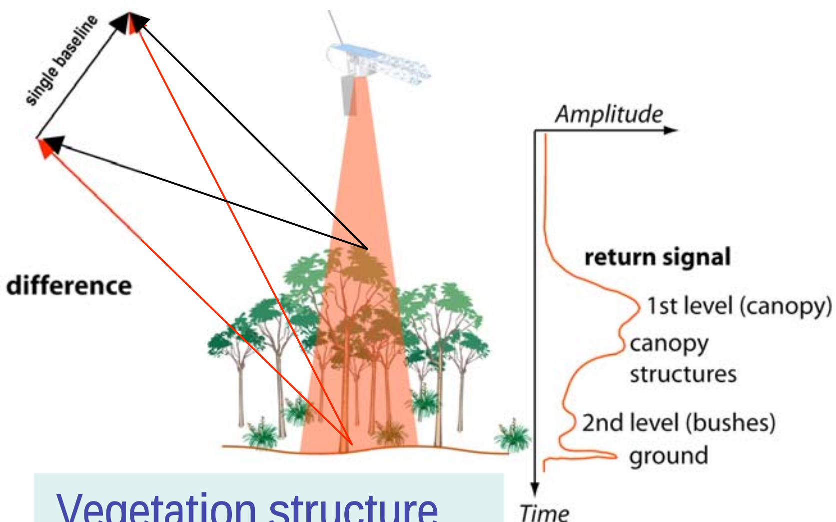

Vegetation structure

Mission Goals and Objectives

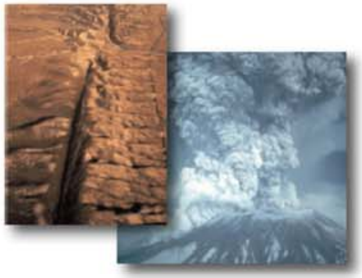

Determine the likelihood of earthquakes, volcanic eruptions, and landslides

US annualized losses from earthquakes are $4.4B/yr yet current hazard maps have an outlook of 30–50 years over hundreds of square kilometers.

Characterize the effects of changing climate and land use on species habitats and carbon budget

The rate of increase [of atmospheric CO_2] over the past century is unprecedented, at least during the past 20,000 years. The structure of ecosystems is a key feature that enables quantification of carbon storage.

Predict the response of ice masses to climate change and impact on sea level

[Ice sheets and glaciers] are exhibiting dramatic changes that are of significant concern for science and international policy. These indicators of climate remain one of the most under-sampled domains in the system.

Application

Understand the behavior of subsurface reservoirs

Management of our hydrological resources is applicable to every state in the union.

QuakeSim Focus is on Deformation Objectives for DESDynI

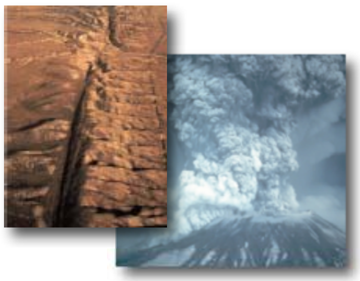

Determine the likelihood of earthquakes, volcanic eruptions, and landslides

US annualized losses from earthquakes are $4.4B/yr yet current hazard maps have an outlook of 30–50 years over hundreds of square kilometers.

- Deformation objectives for solid Earth and cryosphere must be balanced against ecosystem structure objectives in the mission design
  - *Orbit repeat interval*
  - *Coverage*
  - *Radar modes*
- High data volumes for DESDynI drive new paradigm for data processing
- Increase science output by having computational infrastructure in place for science data analysis and interpretation
- Need to balance need for systematic science and response to events such as earthquakes

Deformation Science Requirements

Mission Goals	Science Objectives	Observations	Measurements
Determine the likelihood of earthquakes, volcanic eruptions, and landslides and quantify the magnitude of events	Characterize the nature of deformation at plate boundaries and the implications for earthquake hazards	Measure surface deformation	Coverage of globally actively deforming areas
			3-dimensional (vector)
			100 m imagery
			Accurate to 5% of the rate of the deforming zone with a minimum of 1 mm/yr
			Unaliased temporal sampling, with a minimum of week-timescale measurements immediately following an event
			200 km width imagery across the deforming boundary
		Measure surface disruption	20 m resolution imagery
			400 m zone across the fault
	Characterize how magmatic systems evolve to understand under what conditions volcanoes erupt	Measure surface deformation	Global coverage of Earth's volcanoes
			3-dimensional (vector)
			100 m resolution imagery
			Changes to 1 cm
			Imagery across the area of the volcano
		Measure surface disruption	20 m resolution imagery
			Throughout the area of eruption
	Characterize landslides and detect preslip	Measure surface deformation and disruption	Coverage of landslide prone regions in North America
			3-dimensional (vector)
			Changes to 1 cm
			20 m resolution imagery
			Imagery across the area of the feature

QuakeSim relevance

Similar traceability matrices exist for ecosytems and deformation of ice

Application of QuakeSim to DESDynI Mission Design

- Measurement Requirements
  - *Validate that DESDynI will meet the science objectives*
  - *Assess the quality of DESDynI science products*
  - *Understand observation noise and how it propagates to the science products*
- Deformation baseline models
  - *Construct deformation models and study sensitivity of*
    - Secular fault motions
    - Aquifer subsidence
    - Earthquake displacements
    - Transient motions
  - *From deformation models construct*
    - Synthetic interferograms
    - Time series
  - *Add atmospheric and other noise*
  - *Invert synthetic data*

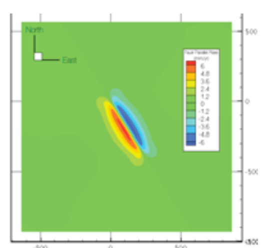

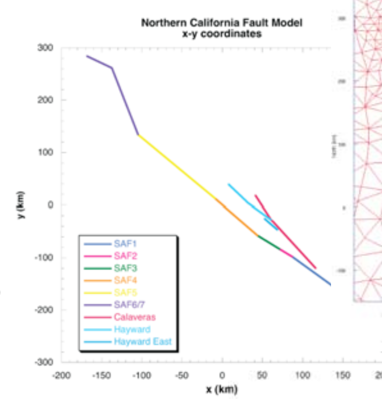

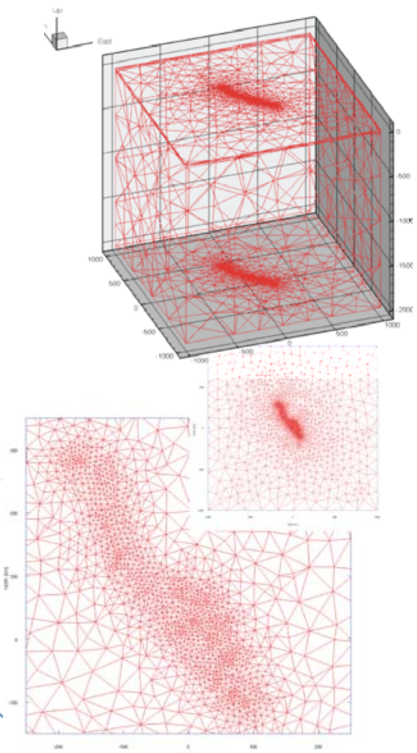

Repeat Interval

- Can discriminate mechanisms for damaging events after 2 years for up to a 14-day repeat interval
  - *Rapid response (6 months) drives need for 8-day orbit*
  - *Ice sheet grounding line studies require repeat be out of phase with the tides driving a maximum allowable repeat of 12 days*
  - *Where available GPS will provide temporal constraints*
- Must balance with coverage needs for the Lidar instrument
  - *Off-pointing to fill in coverage gaps*
  - *Or separate Lidar platform for longer repeat*

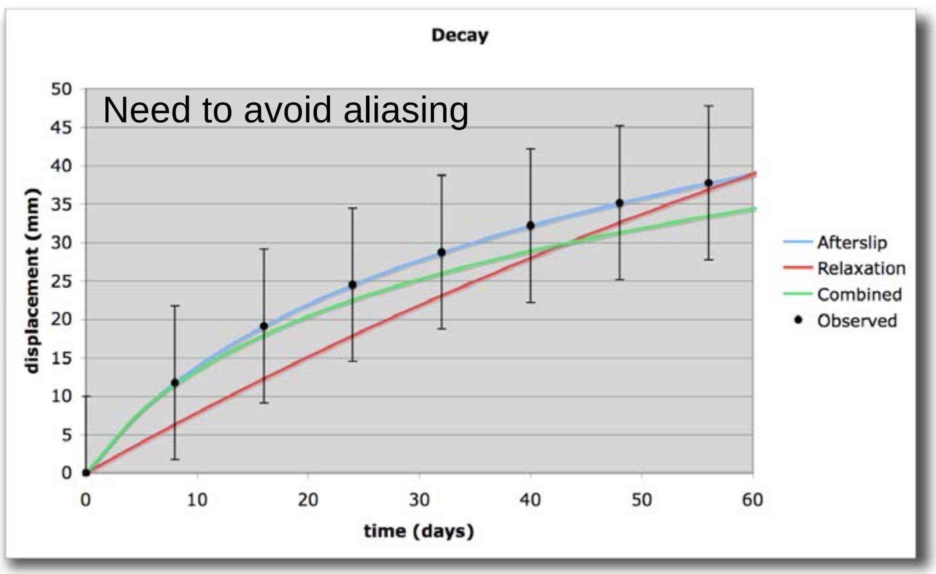

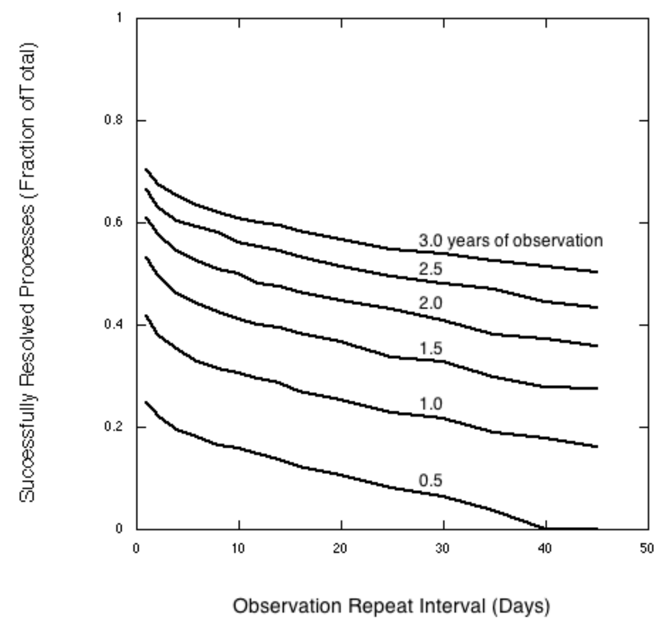

Need for 3D Vector Deformation

- One look provides line-of-sight deformation
  - *Non-unique inversion for fault parameters*
- Need to combine ascending/descending and right/left looks
- Can use QuakeSim for various fault models
  - *Propagate errors for sensitivity analysis*

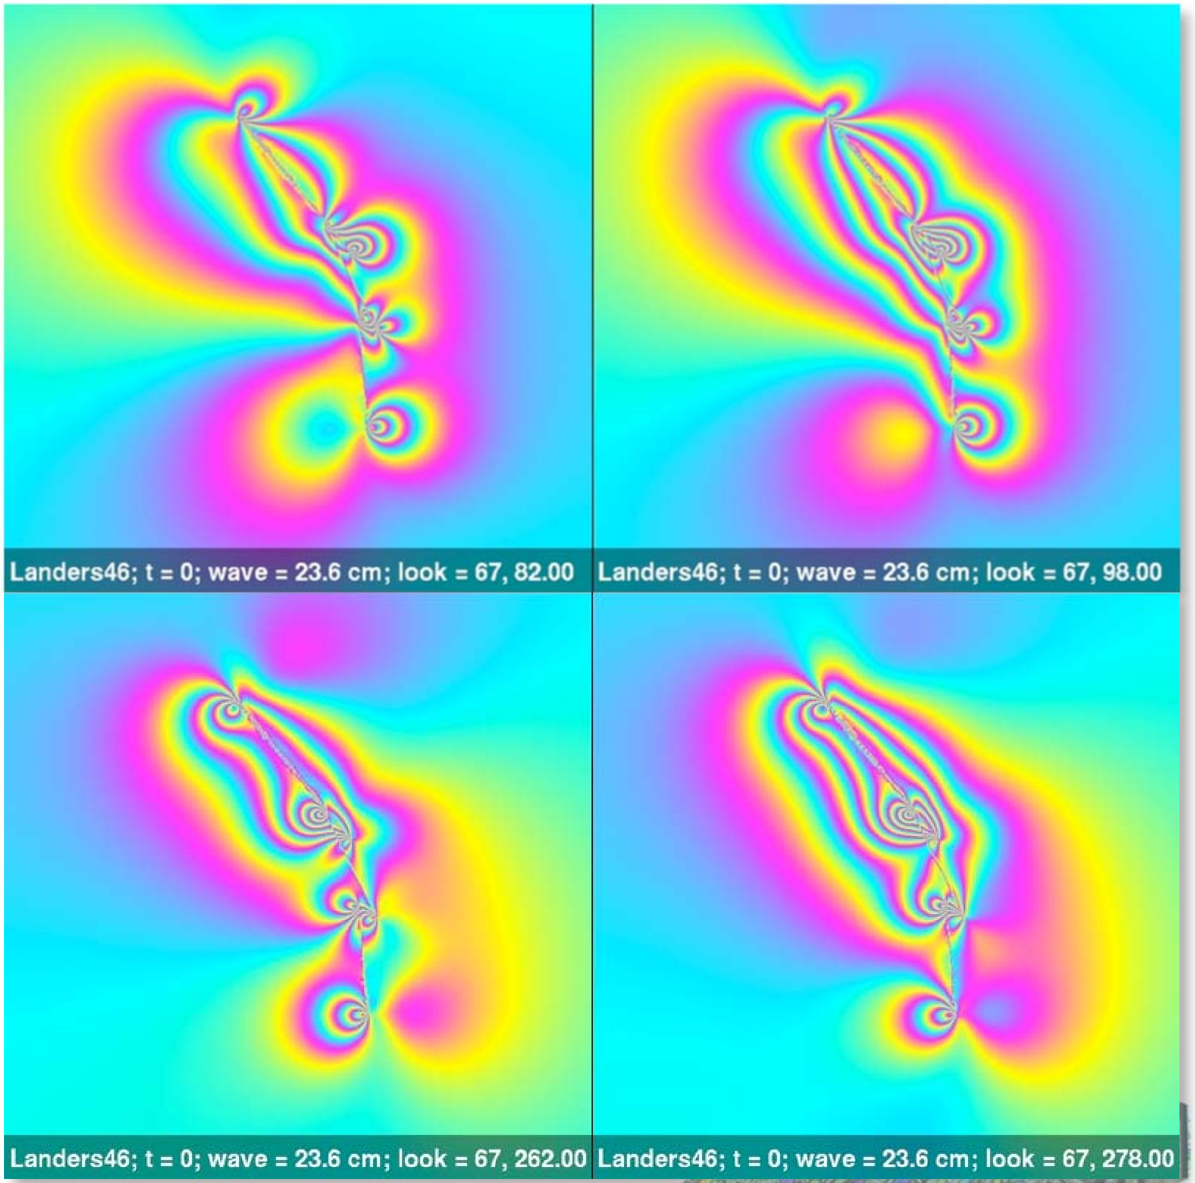

Data Volume for DESDynI

- DESDynI data volume is an order of magnitude more than existing and planned missions
- Once downlinked data must be moved to processing facility and distributed once processed

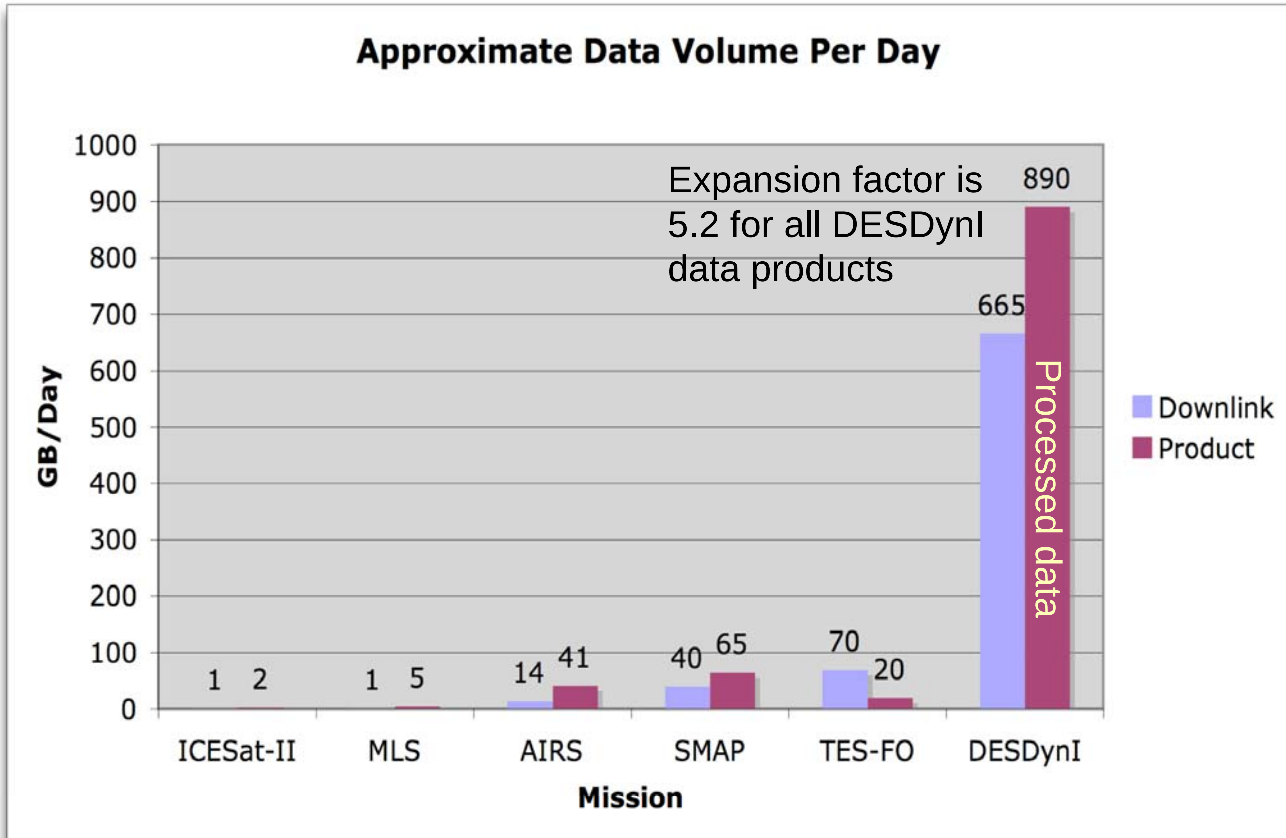

QuakeSim Grid Services Can Serve as Prototype for DESDynI Data Processing

- We are deploying QuakeSim grid services on
  - *Columbia at NASA Ames*
  - *Cosmos at JPL*
  - *NSF TeraGrid*
- For data processing and analysis
  - *Discussing using UAVSAR to prototype the data processing*
  - *GeoFEST is operational*

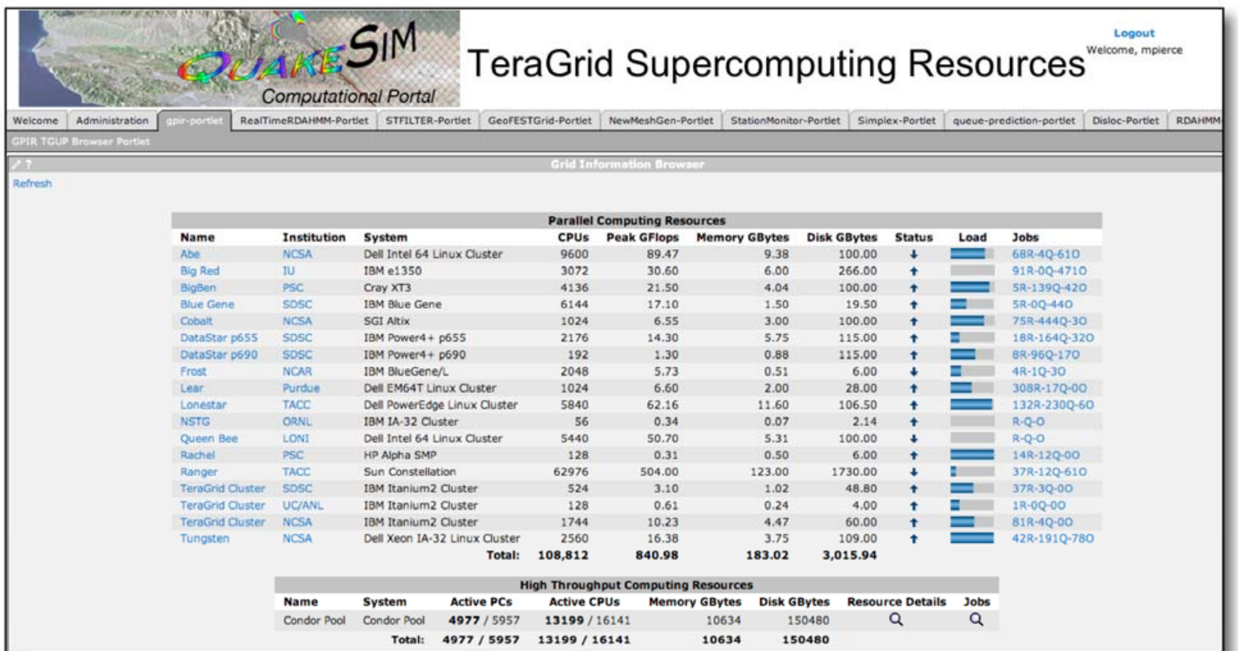

Parallel Computing Resources

Name	Institution	System	CPUs	Peak GFlops	Memory GBytes	Disk GBytes	Status	Load	Jobs
Abe	NCSA	Dell Intel 64 Linux Cluster	9600	89.47	9.38	100.00	↓		68R-4Q-61O
Big Red	IU	IBM e1350	3072	30.60	6.00	266.00	↑		91R-0Q-471O
BigBen	PSC	Cray XT3	4136	21.50	4.04	100.00	↑		5R-139Q-42O
Blue Gene	SDSC	IBM Blue Gene	6144	17.10	1.50	19.50	↑		5R-0Q-44O
Cobalt	NCSA	SGI Altix	1024	6.55	3.00	100.00	↑		75R-444Q-3O
DataStar p655	SDSC	IBM Power4+ p655	2176	14.30	5.75	115.00	↑		18R-164Q-32O
DataStar p690	SDSC	IBM Power4+ p690	192	1.30	0.88	115.00	↑		8R-96Q-17O
Frost	NCAR	IBM BlueGene/L	2048	5.73	0.51	6.00	↓		4R-1Q-3O
Lear	Purdue	Dell EM64T Linux Cluster	1024	6.60	2.00	28.00	↑		308R-17Q-0O
Lonestar	TACC	Dell PowerEdge Linux Cluster	5840	62.16	11.60	106.50	↑		132R-230Q-6O
NSTG	ORNL	IBM IA-32 Cluster	56	0.34	0.07	2.14	↑		R-Q-O
Queen Bee	LONI	Dell Intel 64 Linux Cluster	5440	50.70	5.31	100.00	↓		R-Q-O
Rachel	PSC	HP Alpha SMP	128	0.31	0.50	6.00	↑		14R-12Q-0O
Ranger	TACC	Sun Constellation	62976	504.00	123.00	1730.00	↓		37R-12Q-61O
TeraGrid Cluster	SDSC	IBM Itanium2 Cluster	524	3.10	1.02	48.80	↑		37R-3Q-0O
TeraGrid Cluster	UC/ANL	IBM Itanium2 Cluster	128	0.61	0.24	4.00	↑		1R-0Q-0O
TeraGrid Cluster	NCSA	IBM Itanium2 Cluster	1744	10.23	4.47	60.00	↑		81R-4Q-0O
Tungsten	NCSA	Dell Xeon IA-32 Linux Cluster	2560	16.38	3.75	109.00	↑		42R-191Q-78O
		Total:	**108,812**	**840.98**	**183.02**	**3,015.94**			

High Throughput Computing Resources

Name	System	Active PCs	Active CPUs	Memory GBytes	Disk GBytes	Resource Details	Jobs
Condor Pool	Condor Pool	**4977** / 5957	**13199** / 16141	10634	150480		
	Total:	**4977 / 5957**	**13199 / 16141**	**10634**	**150480**		

DESDynI Deformation Data Analysis

- QuakeSim integrates
  - *InSAR, GPS, fault data*
  - *Inversion and forward models*
  - *Time series analysis and pattern recognizers*
- Will provide computational infrastructure for DESDynI data analysis and interpretation

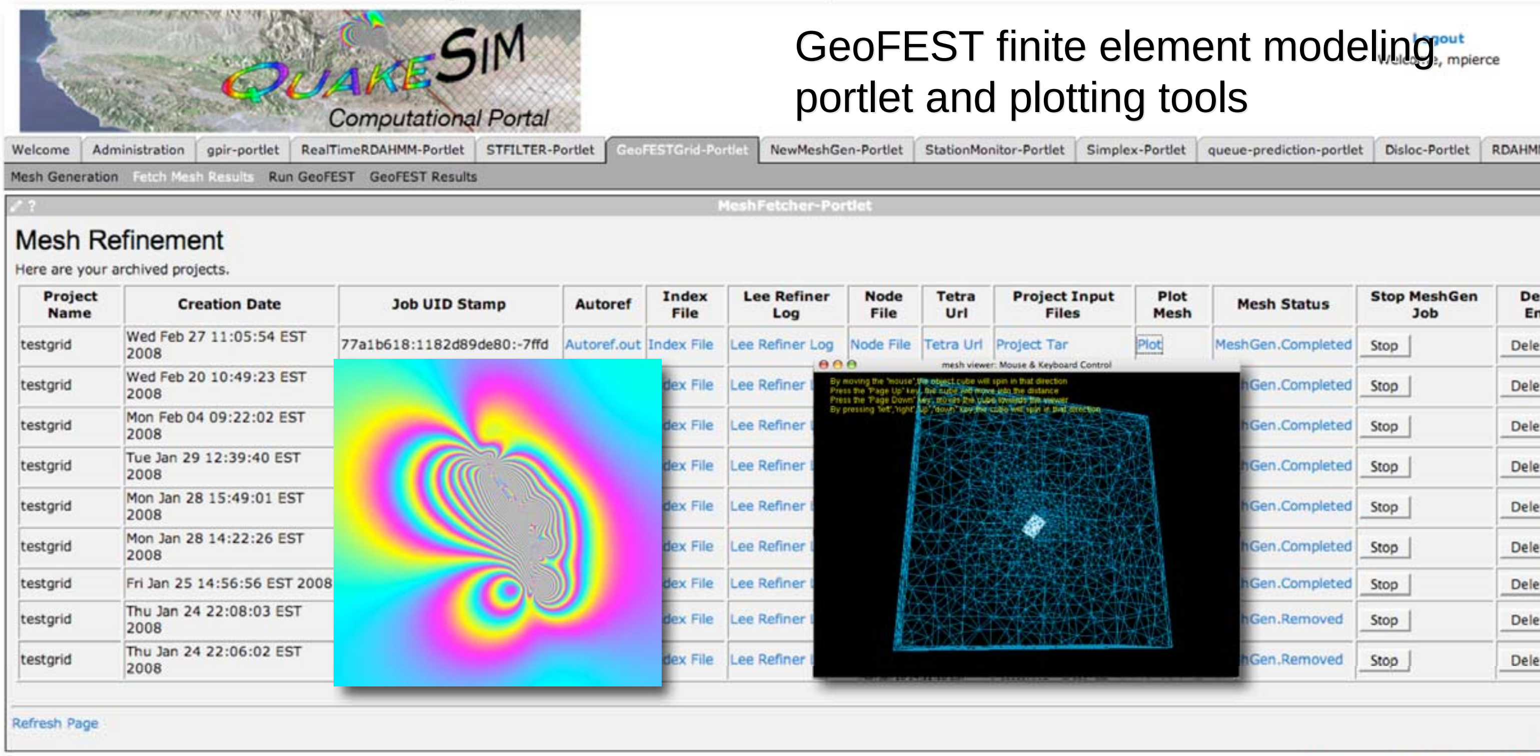

GeoFEST finite element modeling portlet and plotting tools

QuakeSim End-to-End Flow

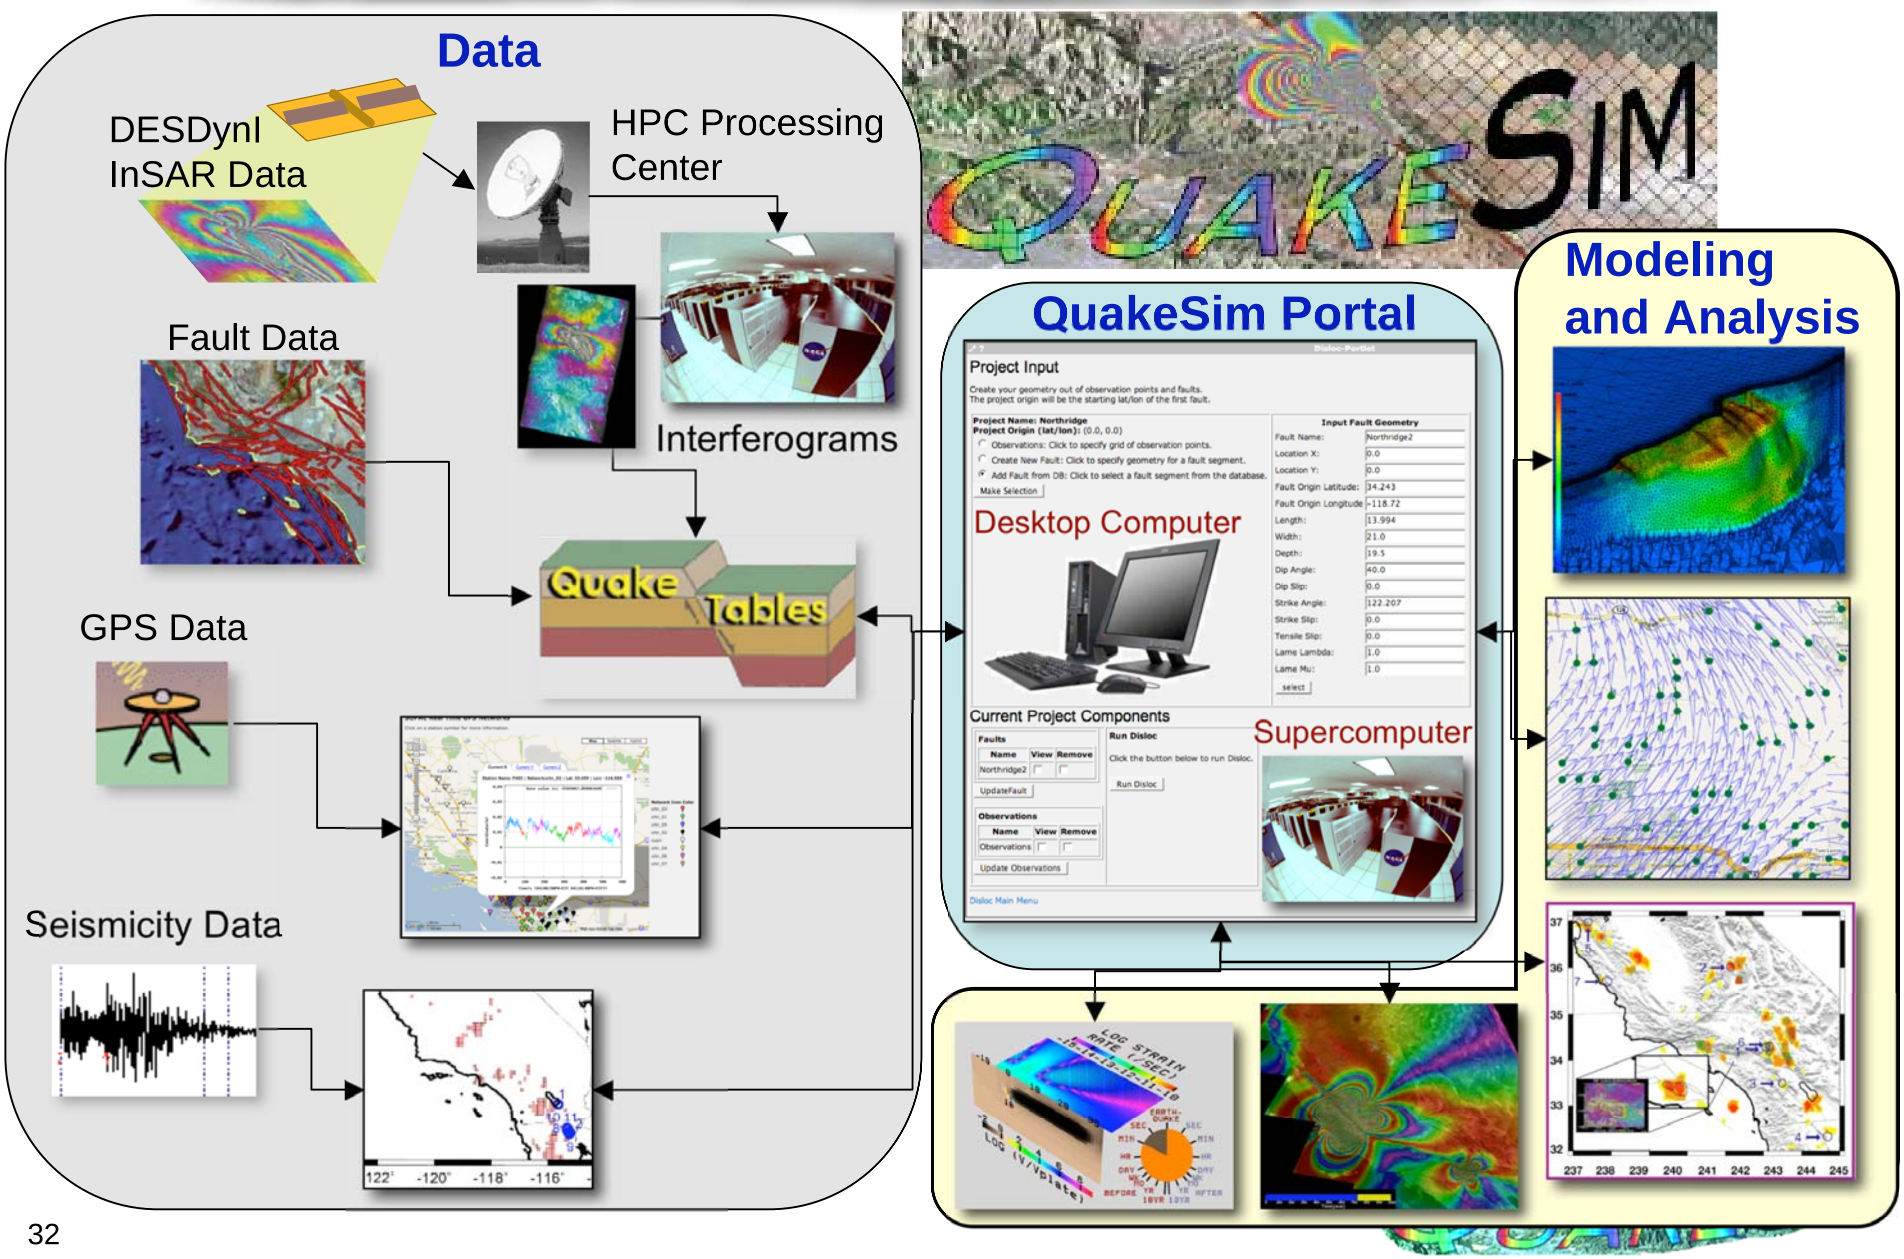

Summary

- Earthquake forecasting methodology is proving extremely successful
- Are analyzing simulated and observed data
  - *Goal is to perform data assimilation of routine data*
- GeoFEST is being used for supporting the crustal deformation InSAR portion of the DESDynI mission
- Developing Grid services with successful implementation on the NSF TeraGrid and plans for JPL and Ames supercomputers
- We have extended the database to include InSAR interferograms
- The portal interface continues to improve